


Search Results

01/15/2002

Search Results for: **[(compiler-controlled <near> bypass*)]**
 Found 4 of 87,654 searched. → Rerun within the Portal


Search within Results


[> Search Help/Tips](#)

[> Advanced Search](#)


Sort by: Title Publication Publication Date Score  Binder


Results 1 - 4 of 4 short listing

- 1 Optimizing communication in HPF programs on fine-grain distributed shared memory 37%

 Satish Chandra , James R. Larus
ACM SIGPLAN Notices , Proceedings of the sixth ACM SIGPLAN symposium on Principles & practice of parallel programming June 1997
 Volume 32 Issue 7
- 2 An evaluation of bottom-up and top-down thread generation techniques 3%

 A. P. W. Böhm , W. A. Najjar , B. Shankar , L. Roh
Proceedings of the 26th annual international symposium on Microarchitecture December 1993
- 3 Lx 2%

 Paolo Faraboschi , Geoffrey Brown , Joseph A. Fisher , Giuseppe Desoli , Fred Homewood
ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture May 2000
 Volume 28 Issue 2
 Lx is a scalable and customizable VLIW processor technology platform designed by Hewlett-Packard and STMicroelectronics that allows variations in instruction issue width, the number and capabilities of structures and the processor instruction set. For Lx we developed the architecture and software from the beginning to support both scalability (variable numbers of identical processing resources) and customizability (special purpose resources). In this paper we consider the followi ...
- 4 Dynamic management of scratch-pad memory space 0%

 M. Kandemir , J. Ramanujam , J. Irwin , N. Vijaykrishnan , I. Kadayif , A. Parikh
Proceedings of the 38th Conference on Design Automation Conference June 2001
 Optimizations aimed at improving the efficiency of on-chip memories are extremely important. We propose a compiler-controlled dynamic on-chip scratch-pad memory (SPM) management framework that uses both loop and data transformations. Experimental results obtained using a generic cost model indicate significant reductions in data transfer activity between SPM and off-chip memory.

Results 1 - 4 of 4 short listing

Search Results

Search Results for: **[(compiler-controlled <near> bypass*)]**
 Found 4 of 87,654 searched. → Rerun within the Portal

Search within Results

> Search Help/Tips



> Advanced Search

Sort by: Title Publication Publication Date Score Binder

Results 1 - 4 of 4 short listing

- 1 Optimizing communication in HPF programs on fine-grain distributed shared memory 37%

Satish Chandra , James R. Larus
 ACM SIGPLAN Notices , Proceedings of the sixth ACM SIGPLAN symposium on Principles & practice of parallel programming June 1997
 Volume 32 Issue 7
- 2 An evaluation of bottom-up and top-down thread generation techniques 3%

A. P. W. Böhm , W. A. Najjar , B. Shankar , L. Roh
 Proceedings of the 26th annual international symposium on Microarchitecture December 1993
- 3 Lx 2%

Paolo Faraboschi , Geoffrey Brown , Joseph A. Fisher , Giuseppe Desoli , Fred Homewood
 ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture May 2000
 Volume 28 Issue 2

Lx is a scalable and customizable VLIW processor technology platform designed by Hewlett-Packard and STMicroelectronics that allows variations in instruction issue width, the number and capabilities of structures and the processor instruction set. For Lx we developed the architecture and software from the beginning to support both scalability (variable numbers of identical processing resources) and customizability (special purpose resources). In this paper we consider the followi ...
- 4 Dynamic management of scratch-pad memory space 0%

M. Kandemir , J. Ramanujam , J. Irwin , N. Vijaykrishnan , I. Kadayif , A. Parikh
 Proceedings of the 38th Conference on Design Automation Conference June 2001

Optimizations aimed at improving the efficiency of on-chip memories are extremely important. We propose a compiler-controlled dynamic on-chip scratch-pad memory (SPM) management framework that uses both loop and data transformations. Experimental results obtained using a generic cost model indicate significant reductions in data transfer activity between SPM and off-chip memory.

Results 1 - 4 of 4 short listing